The Data Center Journal
Where IT, Facilities & Design Meet

Tools
- Home
- Industry Links
- Telecom & Network Consultants
- Data Center Designers
- IT Vendor Listings
- Facilities Vendor Listings
- Data Center Real Estate
- Data Center Equipment For Sale/ Lease
- Help Wanted
- Privacy Policy
- About Us
- Editorial
- Contact Us
- Advertise
- Education
- White Papers
- Glossary

- IT News
- Facilities News
- Design News
- Data Center Journal News
- Press Release News
- Help Wanted News

Evaluation of Direct Current Distribution in Data Centers to Improve Energy Efficiency

By Annabelle Pratt & Pavan Kumar

Data centers are major power consumers. A typical new data center in the United States today has 1,000 racks, occupies 30,00 requires 10 million watts (MW) of power to support the computing infrastructure. Projections for future data centers suggest ever consumption. Future 100,000-square-foot data centers could require 50 MW for the computing infrastructure and an additional 20 MW for c for such a data center would cost nearly $44 million a year (at $0.10 per kilowatt-hour) for services and $18 million annually for c

Intel's most recent efforts to reduce data center power consumption include improving the energy efficiency of server processors multi-core processors and the development of Intel® Core™ microarchitecture. Addressing power consumption at the compute le since every compute watt reduced has a multiplier effect of three. For every watt saved in computation, two additional watts are power conversion and one watt in cooling (the result of no longer having to cool two watts in computation and power conversion). power savings can be gained from small percentages in processor power savings. To add to these power savings, Intel is now loo for ways to spur power savings in other areas of the data center.

Improving Power Conversion Efficiency to Save Power
One potential area for improving energy efficiency that is particularly promising is power conversion. With typical power conversio percent, there is lots of room for improvement. In fact, power conversion efficiency can be increased to above 60 percent simply efficiency components. For example, the typical server power supply unit (one element in a data center's power conversion chain) is only efficient, yet units can be purchased that are 90 percent efficient and pay for their additional cost in less than a year. Adoption of supplies has been slow since server purchases are typically done by IT managers as opposed to data center facility managers. IT focus on getting the best purchase price, whereas a facility manager would look at the long-term power cost as well. As a result, can have a big effect in the long run end up overlooked.

To find ways to improve efficiency throughout the power conversion chain, we evaluated different power delivery architectures. A finding the most efficient power architecture for operating at the wide range of loads experienced by the typical IT data center (a consistently high utilization experienced by high performance computing data centers). At light loads (zero to 25 percent platform utilizatior centers typically burn more energy for power conversion and cooling than for the computer systems themselves.

The Three Power Architectures Considered
We evaluated three different power architectures for large data centers. Each is described in Figures 1 – 3. The percentages given in the c bottom of each power handling component express the power conversion efficiency of that stage, and their product results in the total sy efficiency. These efficiencies are for a heavy workload. Efficiency at a light load can be considerably lower, though for comparative purpos sufficient to consider heavy load only. Light load efficiencies will be addressed in a later section.

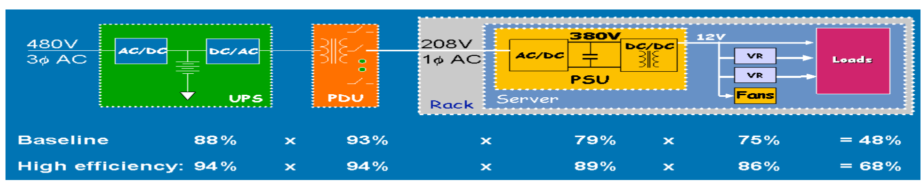

Figure 1: Conventional AC Architecture. In this setup, the utility provides 480 volts to the building and the power delivery train starts with a conversion in a centralized uninterruptible power supply (UPS) which supplies power to many racks. Power is converted to DC to feed an in backup storage system, and then it's converted back to AC and sent to the power distribution unit (PDU). At the PDU voltage is stepped d volts to feed each server in the rack. The PDU efficiency also includes estimated losses in cabling within the data center. The power supply each server converts the 208 volts AC to 380 volts DC. The 380 volts DC is then converted with a DC-to-DC converter to 12 volts DC. Son as hard drives, can take 12 volts directly. Other loads, such as processors, need voltage regulators (VR) to step the voltage down. While t power conversion efficiency for the conventional architecture is just below 50 percent, a system using best-in-class components can achie efficiency. The conventional AC architecture efficiency may be increased further by using a UPS which avoids double conversion. Examples line-interactive UPS and the delta-conversion online UPS, which provide about 98 percent efficiency for the UPS, resulting in 71 percent sys

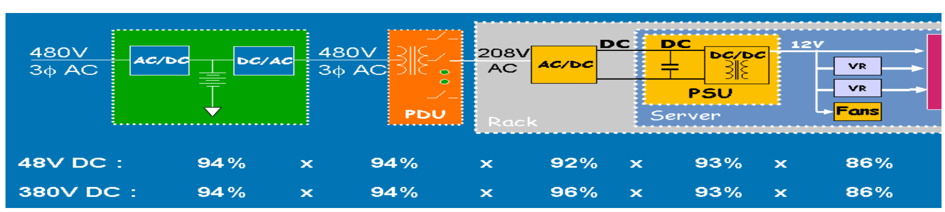

Figure 2: Rack-Level DC Architecture. This setup reduces the heat load and required power supply volume in each server by taking the AC converters out of the power supply for each server and consolidating that function at the rack level with one converter feeding a number has the potential advantage of increased light load efficiency and reduced cooling requirements. By moving the AC-to-DC converters from the rack, the number of converters required to maintain the same level of redundancy is reduced, which reduces initial purchase costs. Th conversion efficiency for the rack-level DC architecture is approximately 65 percent with 48 V DC distribution and 68 percent with 380 V DC is lower with 48 V since the AC/DC conversion stage has to be isolated for 48 V, but a more efficient, non-isolated AC/DC converter may be V DC.

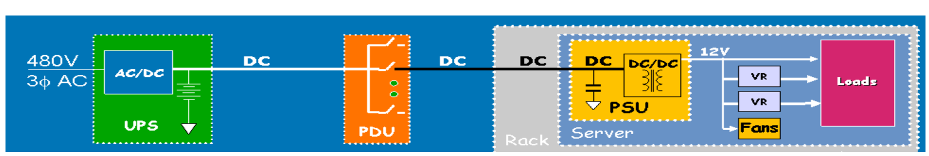

48V DC :	93%	x	97%	x	93%	x	86%	= 72%
380V DC :	97%	x	97%	x	93%	x	86%	= 76%

Figure 3: Facility-Level DC Architecture. Compared to the previous architectures, facility-level DC distribution removes a DC-to-AC conversion stage in the UPS and an AC-to-DC conversion stage in the power supply. There's also no need for a transformer in the PDU. According to our evaluation, facility-level DC distribution improves power conversion efficiency in the data center to about 72 percent at 48 V and 76 percent at 380 V. A limiting factor for the 48 V architecture is that it requires 100 times more copper between the UPS and PDU and about 20 times more copper between the PDU and the rack than an AC architecture requires for similar cable losses. Significant increases in the cost of copper over the past five years make this solution costly for large data centers.

Potential Savings from Facility-level 380 V DC Distribution
We calculated that an efficiency of approximately 75 percent may be achieved with facility-level 380 V DC distribution using best-in-class components. Compared to the typical data center with 50 percent power efficiency and drawing 10 MW for the computing infrastructure load, the 380 V DC data center would draw only 6.67 MW. Assuming a cost of $0.10 per kilowatt (kW) hour, the 380 V DC architecture could save approximately $2.8 million in energy costs per year. Alternatively, the number of servers powered from the same 10 MW power budget could be increased by 60 percent (assuming a cooling system coefficient of performance of 2.6). In addition, DC distribution removes the complications of harmonics and phase balancing, as well as synchronization of parallel UPS connections.2

While AC and 48 V DC distribution are currently used in industry, facility-level 380 V distribution is not. (Rack-level 380 V DC distribution is used though in some high-end servers.) Since 380 V facility-level distribution has been identified as the highest efficiency candidate, Intel, along with several other industry partners, contributed to a small-scale demonstration of 380 V DC facility-level distribution coordinated by Lawrence Berkeley National Laboratory.3 Seven percent input power savings were achieved compared to distribution systems using an AC architecture with best-in-class components. Since most components in a data center operate under light load conditions, measurements were taken with the UPSs loaded at around 35 percent and the server power supplies loaded at about 40 percent.

Figure 4: DC Data Center Demonstration. Photo of the small scale demonstration setup coordinated by Lawrence Berkeley National Laboratory to compare conventional AC architecture (on right) with 380 V DC facility-level distribution (on left).

Challenges in Transitioning to Facility-level DC Distribution Data Centers
Both the 48 V and 380 V facility-level architecture offer significant savings over the conventional AC power architecture and rack-level DC architectures. As mentioned above, the 48 V architecture requires significantly more copper for wiring due to higher currents at the same power levels. It therefore favors a distributed UPS architecture where each data processing cluster has a dedicated DC UPS system feeding it because distributed UPS architectures reduce the amount of wiring required at 48 V. Preliminary analysis suggests the break-even point for such data centers is approximately 5,000 square feet. For smaller data centers, economy of scale will require careful analysis, based on the cost of distributed UPSs, before switching to the new architecture.

For the 380 V DC data center, challenges include establishing standards for wiring, connection and fault protection, and building the industry support needed for volume penetration of this new technology.

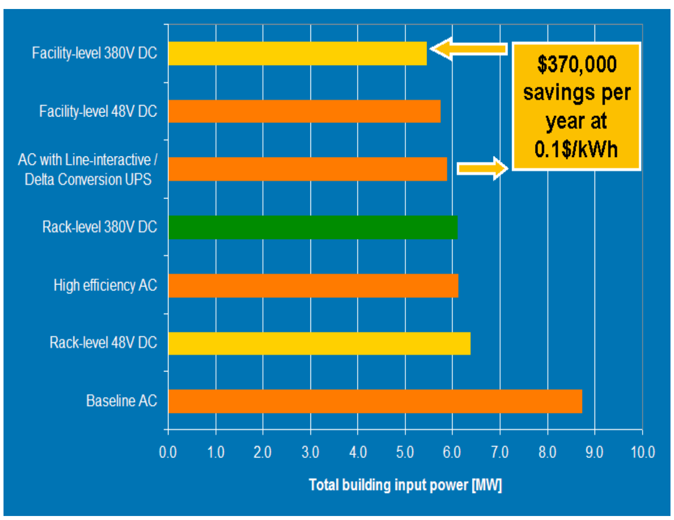

Figure 5: Comparing Input Power Requirements. This chart shows how much money a 380 V DC architecture could save annually over the best-in-class AC architecture in a facility running 10,000 servers with a 300 W load power per server and a 2.6 cooling system coefficient of performance.

Improving Efficiency over a Wide Load Range
The de-rating of power conversion equipment and the redundant nature (due to high availability requirements) of power delivery in data centers compel components in the power delivery train to operate at loads significantly lower than their rating. Because conversion efficiency suffers at these low loads, Intel has been investigating how to improve efficiency over a wide load range.

One source of inspiration comes from Intel's mobile platforms. Intel® Centrino mobile technology-based laptops use multi-phase VRs in which the minimum number of phases required to support the load are used. We investigated applying this technique to power supplies by using multiple small capacity DC/DC converters in each power supply to optimize conversion efficiency to the load condition. Our tests show this could be a particularly effective solution for improving efficiency when loads go below 20 percent of maximum. In the prototype Intel developed, loads at 50 to 100 percent used three DC/DC converters in parallel. For loads from 20 to 50 percent, only two converters were used. For loads of 20 percent and less, just one converter was used. The net effect was a system that could adjust to the load and eliminate the losses normally associated with running at low loads. (See Figure 6.)

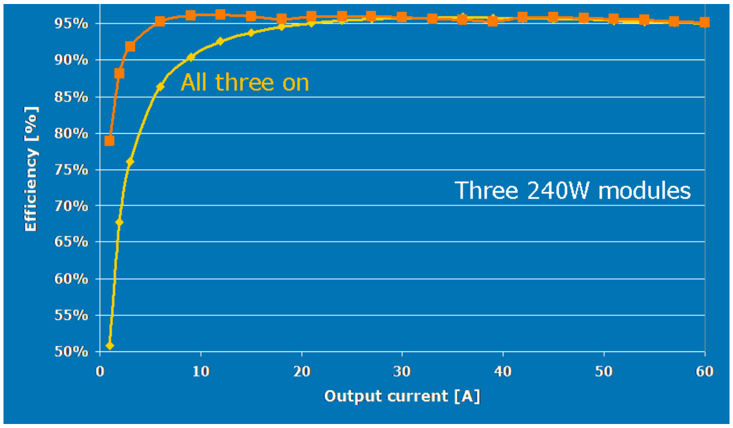

Figure 6: Modular DC/DC Efficiency. Intel experimental data shows that by using three DC/DC converter modules, power supply efficiency can be improved during periods of light loads. At the 2 percent load level, for instance, efficiency improves from 50 percent to 80 percent.

A similar approach could also be used for the UPS stage by using several smaller UPSs in parallel to provide redundancy (an n+1 redundant system). This is more common for 48 V DC UPSs because they are easier to parallel than AC UPSs. It would also work well for 380 V DC output UPSs. This level of modularity would allow IT managers to add hardware as their load grows, reducing initial capital investment.

Power Efficiency in the Data Center Requires Many Solutions

Obviously, the solution to improving the power efficiency of data centers isn't just one solution, but several working in concert. The biggest savings start with improving processor performance per watt because of the multiplier effect it has in reducing the need for power conversion and cooling. Intel's multi-pronged approach has design teams all over the world working in parallel to develop new microarchitectures about every two to three years that will take advantage of the latest process technology advancements. Over the next decade we will be closing in on our goal to enable another tenfold improvement in energy-efficient performance—an achievement that will further increase computing responsiveness and reduce worldwide power consumption for computing devices.

As discussed in this article, data centers can achieve significant savings through using high efficiency power delivery components and a switch to DC power distribution (380 V DC with a central UPS or 48 V using distributed UPSs). This switch is achievable because DC-powered server equipment exists in the same form factor or can readily be built from existing components. In addition, DC system reliability is generally recognized as being as good as, or better than, AC system reliability. Through the comparison of a conventional AC power delivery system to a 380 V DC facility-level system in collaboration with Lawrence Berkeley National Laboratory, we have shown that the efficiency gains from using high efficiency components and the elimination of multiple conversion steps can improve power conversion efficiency in data centers from 50 to 75 percent.

While facility-level distribution at 380 V DC provides the highest system efficiency, a major change in the infrastructure of a data center is required to implement it. In general, a change in mindset from long established design practices will be essential to drive such large infrastructure changes. In this regard, facility-level 380 V DC distribution is primarily proposed for any new data center that is being built. Recognizing the fact that an existing data center cannot change instantaneously, system efficiency can be improved by a series of intermediate steps. The first of this could be in the form of rack-level 380 V DC distribution where AC is supplied to the rack and 380 V DC is distributed within the rack, thus establishing an intermediate architecture. Single racks can be replaced during a natural upgrade cycle without affecting data center infrastructure. Although initially the overall efficiency cannot be dramatically improved, system efficiency can be increased gradually by eliminating PDUs at the end of each aisle while replacing the double conversion UPSs feeding them. Thus over a period of time, an existing AC data center can be transitioned into a high efficiency DC data center.

About the Authors:

Annabelle Pratt is a power research engineer working in the Systems Technology Lab within Intel's Corporate Technology Group. She has been with Intel since 2004 and has focused on technologies and architectures to improve the efficiency of power delivery in servers and data centers.

Prior to Intel, Annabelle developed power supplies for the semiconductor manufacturing and architecture glass-coating industries. She is a senior member of the IEEE and a graduate of Oregon State University.

Pavan Kumar is a power research engineer with Intel's Corporate Technology Group. He focuses on developing power conversion technologies for Intel's future platforms. He has been with Intel since 1999. His areas of interest in power delivery range from servers to handhelds. Prior to joining Intel he was involved in developing power converter solutions for telecommunication equipment with Nortel Networks, Ottawa, Canada.

The authors would like to acknowledge Tomm V. Aldridge, Intel Principal Engineer and CTG Power Architect, and Kevin Bross, Modular Systems Architect, for their contributions.

1 C. Patel, et. al, "Thermal Considerations in Cooling Large Scale High Compute Density Data Centers," ITHERM 2002, pp. 767 – 776, 2002.

2 Marquet, D., et al., "New Flexible Powering Architecture for Integrated Service Operators", Intelec 2005 Conf. Proc., pp. 575-80, 2005.

3 See http://hightech.lbl.gov/dc-powering/